

Bounding the Probability of Causation in Mediation Analysis

A. P. Dawid*

R. Murtas[†]M. Musio[‡]

November 12, 2014

Abstract

Given empirical evidence for the dependence of an outcome variable on an exposure variable, we can typically only provide bounds for the “probability of causation” in the case of an individual who has developed the outcome after being exposed. We show how these bounds can be adapted or improved if further information becomes available. In addition to reviewing existing work on this topic, we provide a new analysis for the case where a mediating variable can be observed. In particular we show how the probability of causation can be bounded when there is no direct effect and no confounding.

Keywords: Causal inference, Mediation Analysis, Probability of Causation

1 Introduction

Many statistical analyses aim at a causal explanation of the data. In particular in epidemiology many studies are conducted to try to understand if and when an exposure will cause a particular disease. Also in a Court of Law, when we want to assess legal responsibility we usually refer to causality. But when discussing this topic it is important to specify the exact query we want to talk about. For example it may be claimed in court that it was Ann’s taking the drug that was the cause of her death. This type of question relates to the cause of an observed effect (“CoE”) and is fundamental to the allocation of responsibility. On the other hand much of classical statistical design and analysis, for example randomized agricultural or medical experiments, has been created to address questions about the effects of applied causes (“EoC”). When we address an EoC query, we are typically asking a hypothetical question: “What would happen to Ann if she were to take the drug?”. At the very same time we can address alternative hypothetical questions: “What would happen to Ann if she were not to take the drug?”.

Assessing the effects of causes can be achieved in straightforward fashion using a framework based on probabilistic prediction and statistical decision theory [2]. To formalize the

*University of Cambridge apd@statslab.cam.ac.uk

[†]University of Cagliari ro.murtas@gmail.it

[‡]University of Cagliari mmusio@unica.it

	Die	Live	Total
Exposed	30	70	100
Unexposed	12	88	100

Table 1: Deaths in individuals exposed and unexposed to the same drug taken by Ann

problem, let X be a binary decision variable denoting whether or not Ann takes the drug, and Y the response, coded as 1 if she dies and 0 if not. We denote by P_1 [resp. P_0] the probability distribution of Y ensuing when X is set to the value 1 [resp. 0]. The two distributions P_1 and P_0 are all that is needed to address EoC-type queries: I can compare these two different hypothetical distributions for Y , decide which one I prefer, and take the associated decision.

The situation is different for a CoE query, where the drug has already been taken and the outcome observed. A natural way to address a CoE question would be to try to imagine what would have happened to Ann had she not been taken the drug. In other words, given the fact that Ann actually took the drug and died, how likely is it that she would not have died if she had not taken the drug? We can not address a CoE query using only the two distribution P_1 and P_0 . In fact we can no longer base our approach purely on the probability distribution of Y and X conditioned on known facts, since we know the values of both variables ($Y = 1, X = 1$), and after conditioning on that knowledge there is no probabilistic uncertainty left to work with. Nevertheless we want an answer. This query can be approached by introducing (for any individual) an associated pair of “potential responses” $\mathbf{Y} := (Y(0), Y(1))$, where $Y(x)$ denotes the value of the response Y that will be realized when the exposure X is set to x (which we write as $X \leftarrow x$). Both potential responses are regarded as existing, simultaneously, prior to the choice of X , the actual response Y then being determined as $Y = Y(X)$. However, for any each individual just one of the potential responses will be observable. For example, only $Y(1)$ will be observable if in fact $X \leftarrow 1$; $Y(0)$ will then be *counterfactual*, because it relates to a situation, $X \leftarrow 0$, which is contrary to the known fact $X \leftarrow 1$.

To address the court’s query we use the formulation of *Probability of Causation*, PC, as given by Pearl in [5] (where it is named *Probability of Necessity*). In terms of the triple $(X_A, Y_A(0), Y_A(1))$, we define the Probability of Causation in Ann’s case as:

$$PC_A = P_A(Y_A(0) = 0 \mid X_A = 1, Y_A(1) = 1) \quad (1)$$

where P_A denotes the probability distribution over attributes of Ann. Knowing that Ann did take the drug ($X_A = 1$) and the actual response was recovery ($Y_A = 1$), this is the probability that the potential response $Y_A(0)$, that would been observed had Ann not taken the drug, would have been different ($Y_A(0) = 0$). But how are we to get a purchase on this quantity?

Suppose that a good experimental study tested the same drug taken by Ann, and produced the data reported in Table 1. Since our analysis here is not concerned with purely statistical variation due to small sample sizes, we take proportions computed from this table as accurate estimates of the corresponding population probabilities (see [3] for issues related to the use of small-sample data for making causal inferences). Thus we take

$$\begin{aligned} \Pr(Y = 1 \mid X \leftarrow 1) &= 0.30 \\ \Pr(Y = 1 \mid X \leftarrow 0) &= 0.12 \end{aligned}$$

where we use \Pr to denote population probabilities.

We see that, in the experimental population, individuals exposed to the drug ($X \leftarrow 1$) were more likely to die than those unexposed ($X \leftarrow 0$), by 28 percentage points. So can the court infer that was Ann’s taking the drug that caused her death? More generally: Is it correct to use such experimental results, concerning a population, to say something about a single individual? This “Group-to-individual” (G2i) issue is discussed by Dawid *et al.* [6] in relation to the question “When can Science be relied upon to answer factual disputes in litigation?”. It is there pointed out that in general we can not obtain a point estimate for PC_A ; but we can provide useful information, in the form of bounds between which this quantity must lie.

In this paper we show how these bounds can be adapted or improved when further information is available. In §2 we consider the basic situation where we have information only on exposure and outcome. In §3 we bound the probability of causation when we have additional information on a pre-treatment covariate. Section 4 considers the situation in which unobserved variables confound the exposure-outcome relationship. Finally in §5 we introduce new bounds for PC when a mediating variable can be observed. Section 6 presents some concluding comments.

2 Starting Point: Simple Analysis

In this section we discuss the simple situation in which we have information, as in Table 1, from a randomized experimental study. We need to assume that the fact of Ann’s exposure, X_A , is independent of her potential responses \mathbf{Y}_A :

$$X_A \perp\!\!\!\perp \mathbf{Y}_A. \quad (2)$$

Property (2) parallels the “no-confounding” property $X_i \perp\!\!\!\perp \mathbf{Y}_i$ which holds for individuals i in the experimental study on account of randomization. We further suppose that Ann is exchangeable with the individuals in the experiment, *i.e.* she could be considered as a subject in the experimental population.

On account of (2) and exchangeability, (1) reduces to $PC_A = \Pr(Y(0) = 0 \mid Y(1) = 1)$; but we can not fully identify this from the data. In fact we can never observe the joint event $(Y(0) = 0; Y(1) = 1)$, since at least one of $Y(0), Y(1)$ must be counterfactual. In particular, we can never learn anything about the dependence between $Y(0)$ and $Y(1)$. However, even without making any assumptions about this dependence, we can derive the following inequalities [3]:

$$1 - \frac{1}{RR} \leq PC_A \leq \frac{\Pr(Y = 0 \mid X \leftarrow 0)}{\Pr(Y = 1 \mid X \leftarrow 1)} \quad (3)$$

where

$$RR = \frac{\Pr(Y = 1 \mid X \leftarrow 1)}{\Pr(Y = 1 \mid X \leftarrow 0)}$$

is the *experimental risk ratio* between exposed and unexposed. And these bounds can be estimated from the experimental data using the population death rates computed in §1.

In many cases of interest (such as Table 1), we will have

$$\Pr(Y = 1 \mid X \leftarrow 0) < \Pr(Y = 1 \mid X \leftarrow 1) < \Pr(Y = 0 \mid X \leftarrow 0).$$

Then the lower bound in (3) will be non-trivial, while the upper bound will exceed 1, and hence be vacuous.

We see from (3) that whenever $RR > 2$ the Probability of Causation PC_A will exceed 50%. In a civil court this is often taken as the criterion to assess legal responsibility “on the balance of probabilities” (although the converse is false: it would not be correct to infer $PC_A < .5$ from the finding $RR < 2$). Since, in Table 1, the exposed are 2.5 times as likely to die as the unexposed ($RR = 30/12 = 2.5$), we have enough confidence to infer causality in Ann’s case: We have $0.60 \leq PC_A \leq 1$.

3 Additional Covariate Information

In this Section we show how we can refine the bounds of (3) if further information about a pre-treatment covariate S is available. We now take the assumptions of § 2 to hold after conditioning on S (indeed in cases where the original assumptions fail, it may well be possible to reinstate them by conditioning on a suitable covariate S). In particular, $X_A \perp\!\!\!\perp Y_A \mid S_A$, and $X_i \perp\!\!\!\perp Y_i \mid S_i$: adjusting for S is enough to control for confounding, both for Ann and in the study.

3.1 Fully observable

Consider first the situation where we can observe S both in the experimental data and in Ann. We can apply the analysis of § 2, after conditioning on S , to obtain the estimable lower bound:

$$1 - \frac{1}{RR(s_A)} \leq PC_A \quad (4)$$

where

$$RR(s) = \frac{\Pr(Y = 1 \mid X \leftarrow 1, S = s)}{\Pr(Y = 1 \mid X \leftarrow 0, S = s)},$$

and s_A is Ann’s value for S .

3.2 Observable in data only

But even when we can only observe S in the population, and not in Ann, we can sometimes refine the bounds in (3). Thus suppose S is binary, and from the data we infer the following probabilities (which are consistent with the data of Table 1):

$$P_A(S = 1) = 0.50 \quad (5)$$

$$P_A(Y = 1 \mid X \leftarrow 1, S = 1) = 0.60 \quad (6)$$

$$P_A(Y = 1 \mid X \leftarrow 0, S = 1) = 0 \quad (7)$$

$$P_A(Y = 1 \mid X \leftarrow 1, S = 0) = 0 \quad (8)$$

$$P_A(Y = 1 \mid X \leftarrow 0, S = 0) = 0.24 \quad (9)$$

Since we know $X_A = 1, Y_A = 1$, from (8) we deduce $S_A = 1$, and so $Y_A(0) = 0$ by (7). That is, in this special case we can infer causation in Ann’s case—even though we have not directly observed her value for S .

	Die	Live	Total
Exposed	18	82	100
Unexposed	24	76	100

Table 2: Non experimental data

More generally (see [1]) we can refine the bounds in (3) as follows:

$$\frac{\Delta}{\Pr(Y = 1 \mid X \leftarrow 1)} \leq \text{PC} \leq 1 - \frac{\Gamma}{\Pr(Y = 1 \mid X \leftarrow 1)} \quad (10)$$

where

$$\begin{aligned} \Delta &= \sum_s \Pr(S = s) \times \max \{0, \Pr(Y = 1 \mid X \leftarrow 1, S = s) - \Pr(Y = 1 \mid X \leftarrow 0, S = s)\} \\ \Gamma &= \sum_s \Pr(S = s) \times \max \{0, \Pr(Y = 1 \mid X \leftarrow 1, S = s) - \Pr(Y = 0 \mid X \leftarrow 0, S = s)\} \end{aligned}$$

These bounds are never wider than those obtained from (3), which ignores S .

4 Unobserved Confounding

So far we have assumed no confounding: $X \perp\!\!\!\perp Y$ (perhaps conditionally on a suitable covariate S), both for Ann and for the study data. Now we drop this assumption for Ann. Then the experimental data can not be used, by themselves, to learn about $\text{PC}_A = \Pr(Y_A(0) = 0 \mid X_A = 1, Y_A(1) = 1)$.

We might however be able to gather additional *observational* data, having the same dependence between X and Y as for Ann. Let Q denote the joint observational distribution of (X, Y) , estimable from such data. Tian and Pearl [4] obtain the following bounds for PC_A , given both experimental and nonexperimental data:

$$\max \left\{ 0, \frac{Q(Y = 1) - \Pr(Y = 1 \mid X \leftarrow 0)}{Q(X = 1, Y = 1)} \right\} \leq \text{PC}_A \leq \min \left\{ 1, \frac{\Pr(Y = 0 \mid X \leftarrow 0) - Q(X = 0, Y = 0)}{Q(X = 1, Y = 1)} \right\}. \quad (11)$$

For example, suppose that, in addition to the data of Table 1, we have observational data as in Table 2.

Thus

$$\begin{aligned} Q(Y = 1) &= 0.21 \\ Q(X = 1, Y = 1) &= 0.09 \\ Q(X = 0, Y = 0) &= 0.38. \end{aligned}$$

Also, from Table 1 we have $\Pr(Y = 1 \mid X \leftarrow 0) = 0.12$ (so $\Pr(Y = 0 \mid X \leftarrow 0) = 1 - 0.12 = 0.88$). From (11) we thus find $1 \leq \text{PC}_A \leq 1$. We deduce that Ann would definitely have survived had she not taken the drug.



Figure 1: Directed Acyclic Graph representing a mediator M , responding to exposure X and affecting response Y . There is no “direct effect”, unmediated by M , of X on Y .

5 Mediation Analysis

In this Section we bound the Probability of Causation for a case where a third variable, M , is involved in the causal pathway between the exposure X and the outcome Y . Such a variable is called a *mediator*. In general, the total causal effect of X on Y can be split into two different effects: One mediated by M (the *indirect effect*) and one not so mediated (the *direct effect*). Here we shall only consider the case of no direct effect, as intuitively described by Figure 1. We shall be interested in the case that M is observable in the experimental data, but is not observed for Ann, and see how this additional experimental evidence can be used to refine the bounds on PC_A .

To formalize our assumption of “no direct effect”, we introduce $M(x)$, the potential value of M for $X \leftarrow x$, and $Y(m)$, the potential value of Y for $M \leftarrow m$ and $X \leftarrow x$, where x is any value—the irrelevance of that value representing the property that X has no effect on Y over and above that transmitted through its influence on the mediator M . The potential value of Y for $X \leftarrow x$ (in cases where there is no intervention on M , which we here assume) is then $Y^*(x) := Y\{M(x)\}$.

In the sequel we restrict to the case that all variables are binary, and define $\mathbf{M} := (M(0), M(1))$, $\mathbf{Y} := (Y(0), Y(1))$, $\mathbf{Y}^* := (Y^*(0), Y^*(1))$. In particular, we have observable variables $(X, M, Y) = (X, M(X), Y^*(X))$. We denote the bivariate distributions of the potential response pairs by

$$\begin{aligned} m_{ab} &:= \Pr(M(0) = a, M(1) = b) \\ y_{rs} &:= \Pr(Y(0) = r, Y(1) = s) \\ y_{rs}^* &:= \Pr(Y^*(0) = r, Y^*(1) = s). \end{aligned}$$

Then

$$\begin{aligned} m_{a+} &= \Pr(M = a \mid X \leftarrow 0) \\ m_{+b} &= \Pr(M = b \mid X \leftarrow 1) \\ y_{r+} &= \Pr(Y = r \mid M \leftarrow 0) \\ y_{+s} &= \Pr(Y = s \mid M \leftarrow 1) \\ y_{r+}^* &= \Pr(Y = r \mid X \leftarrow 0) \\ y_{+s}^* &= \Pr(Y = s \mid X \leftarrow 1) \end{aligned}$$

where y_{r+} denote $\sum_{s=0}^1 y_{rs}$, etc.

In addition to the assumptions of §2 we further suppose that none of the causal mechanisms depicted in Figure 1 are confounded—expressed mathematically by assuming mutual independence between X , \mathbf{M} and \mathbf{Y} (both for experimental individuals, and for Ann). Then m_{a+} , m_{+b} , y_{r+} , y_{+s} , y_{r+}^* , y_{+s}^* are all estimable from experimental data in which X is randomized, and M and Y are observed.

It is also then easy to show the Markov property:

$$Y \perp\!\!\!\perp X \mid M.$$

This observable property can serve as a test of the validity of our conditions. It implies

$$y_{00}^* = m_{00}y_{0+} + (m_{01} + m_{10})y_{00} + m_{11}y_{+0} \quad (12)$$

$$y_{01}^* = m_{01}y_{01} + m_{10}y_{10} \quad (13)$$

$$y_{10}^* = m_{01}y_{10} + m_{10}y_{01} \quad (14)$$

$$y_{11}^* = m_{00}y_{1+} + (m_{01} + m_{10})y_{11} + m_{11}y_{+1}, \quad (15)$$

and

$$y_{r+}^* = m_{0+}y_{r+} + m_{1+}y_{+r} \quad (16)$$

$$y_{+s}^* = m_{+0}y_{s+} + m_{+1}y_{+s} \quad (17)$$

Suppose now that we observe $X_A = 1$, $Y_A = 1$, but do not observe M_A . We have

$$\text{PC}_A = \frac{y_{01}^*}{y_{+1}^*} = \frac{m_{01}y_{01} + m_{10}y_{10}}{y_{+1}^*}. \quad (18)$$

The denominator of (18) is $\Pr(Y = 1 \mid X \leftarrow 1)$, which is estimable from the data.

As for the numerator, this can be expressed as

$$2\mu\eta + A\mu + B\eta + AB = 2(\mu + B/2)(\eta + A/2) + AB/2 \quad (19)$$

with $\mu = m_{01}$, $\eta = y_{01}$, $A = y_{+0} - y_{0+}$, $B = m_{+0} - m_{0+}$. Note that A , B are identified from the data, while for μ and η we can only obtain inequalities:

$$\begin{aligned} \max\{0, -B\} &\leq \mu \leq \min\{m_{0+}, m_{+1}\} \\ \max\{0, -A\} &\leq \eta \leq \min\{y_{0+}, y_{+1}\}, \end{aligned}$$

so that

$$\begin{aligned} |B/2| &\leq \mu + B/2 \leq \min\{\frac{1}{2}(m_{0+} + m_{+0}), \frac{1}{2}(m_{1+} + m_{+1})\} \\ |A/2| &\leq \eta + A/2 \leq \min\{\frac{1}{2}(y_{0+} + y_{+0}), \frac{1}{2}(y_{1+} + y_{+1})\}. \end{aligned} \quad (20)$$

The lower [resp., upper] limit for (19) will be when $\mu + B/2$ and $\eta + A/2$ are both at their lower [resp., upper] limits. In particular, the lower limit for (19) is $\max\{0, AB\}$. Using (16) and (17), we compute $AB = y_{+1}^* - y_{1+}^*$, which leads to the lower bound

$$\text{PC}_A \geq 1 - \frac{\Pr(Y = 1 \mid X \leftarrow 1)}{\Pr(Y = 1 \mid X \leftarrow 0)} = 1 - \frac{1}{\text{RR}}, \quad (21)$$

exactly as for the case that M was not observed. Thus the possibility to observe a mediating variable in the experimental data has not improved our ability to lower bound PC_A .

We do however obtain an improved upper bound. Taking into account the various possible choices for the upper bounds in (20), the upper bound for the numerator of (18), in terms of experimentally estimable quantities, is given in Table 3.

	$m_{1+} + m_{+1} \geq 1$	$m_{1+} + m_{+1} < 1$
$y_{1+} + y_{+1} \geq 1$	$m_{0+}y_{0+} + m_{+0}y_{+0}$	$m_{1+}y_{+0} + m_{+1}y_{0+}$
$y_{1+} + y_{+1} < 1$	$m_{0+}y_{+1} + m_{+0}y_{1+}$	$m_{1+}y_{1+} + m_{+1}y_{+1}$

Table 3: Upper bound for numerator of (18)

5.1 Example

Suppose we obtain the following values from the data:

$$\begin{aligned}
\Pr(M = 1 \mid X \leftarrow 1) &= 0.25 \\
\Pr(M = 1 \mid X \leftarrow 0) &= 0.025 \\
\Pr(Y = 1 \mid M \leftarrow 1) &= 0.9 \\
\Pr(Y = 1 \mid M \leftarrow 0) &= 0.1
\end{aligned}$$

(these are consistent with Table 1).

Then we find $0.60 \leq \text{PC}_A \leq 0.76$; whereas without taking account of the mediator we would have no non-trivial upper bound.

6 Discussion

In this paper we have considered estimation of the Probability of Causation in a number of contexts, including a novel analysis for the case of a mediating variable, in the absence of a direct effect. As we saw in §5 considering such a third variable in the pathway between exposure and outcome can lead to an improved upper bound, although conclusions about the lower bound remains the same.

Even if the case of no direct effect is special and unusual, there certainly do exist cases, such as the relationship between anxiolytics and cars crash mediated by alcohol consumption, or the relationship between aspirin and yellow jaundice due to hemolytic anemia mediated by favism, where this assumption is plausible. The next step will be generalize our analysis to more general cases of mediation, allowing for a direct effect and for unobserved confounding.

References

- [1] Dawid, A. P. (2011). The role of scientific and statistical evidence in assessing causality. In *Perspectives on Causation* (R. Goldberg, Ed.). Oxford: Hart Publishing, 133–147.
- [2] Dawid, A. P. (2014). Statistical causality from a decision-theoretic perspective. *Annual Review of Statistics and Its Application* **2**. In Press.
- [3] Dawid, A. P., Musio, M., and Fienberg, S. E. (2014). From statistical evidence to evidence of causality. [arXiv:1311.7513](https://arxiv.org/abs/1311.7513).
- [4] Tian, J. and Pearl, J. (2000). Probabilities of causation: Bounds and identification. *Annals of Mathematics and Artificial Intelligence* **28**, 287–313.

- [5] Pearl, J. (1999). Probabilities of causation: Three counterfactual interpretations and identification. *Synthese* **121**, 93-149.
- [6] Dawid, A. P., Fienberg, S. and Faigman, D. (2014a). Fitting science into legal contexts: Assessing effects of causes or causes of effects? *Sociological Methods and Research* **43**, 359–390.